
Subject: VLAN Setup

Posted by [jonasb](#) on Thu, 09 Aug 2007 06:07:11 GMT

[View Forum Message](#) <> [Reply to Message](#)

My server are connected to a switch with cisco trunking.
I have 2 VLANs. How can I create VEs in both VLANs?

Subject: Re: VLAN Setup

Posted by [vaverin](#) on Thu, 09 Aug 2007 07:33:17 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Jonas,

I hope the following link will help you:

<http://wiki.openvz.org/VLAN>

Please fell free to ask me if you have some questions.

thank you,

Vasily Averin

Subject: Re: VLAN Setup

Posted by [jonasb](#) on Thu, 09 Aug 2007 08:09:47 GMT

[View Forum Message](#) <> [Reply to Message](#)

Thanks! I have seen that before. Is it possible to not using VLAN inside the VE? Just using vconfig outside the VE.

Subject: Re: VLAN Setup

Posted by [vaverin](#) on Thu, 09 Aug 2007 08:14:59 GMT

[View Forum Message](#) <> [Reply to Message](#)

Of course, just use "vconfig add eth0 <vlan_id>" on HW node.

Subject: Re: VLAN Setup

Posted by [jonasb](#) on Thu, 09 Aug 2007 08:57:39 GMT

[View Forum Message](#) <> [Reply to Message](#)

I can't get it to work. I can start the VE. And I can ping it from HN but not outside HN. The strange thing is that my vlan disappear when I start the VE. And is available again when I stop it.

```
[root@vz-test ~]# uname -a
Linux vz-test 2.6.18-8.1.8.el5.028stab039.1PAE #1 SMP Mon Jul 23 20:45:26 MSD 2007 i686
i686 i386 GNU/Linux
```

```
[root@vz-test ~]# vconfig add eth1 105
Added VLAN with VID == 105 to IF -:eth1:-
```

```
[root@vz-test ~]# ifconfig eth1.105 192.168.20.137 netmask 255.255.255.240 broadcast
192.168.20.143
```

```
[root@vz-test ~]# vzctl set 10666 --netdev_add eth1.105 --save
Saved parameters for VE 10666
```

```
[root@vz-test ~]# vzctl set 10666 --ipadd 192.168.20.138 --save
Saved parameters for VE 10666
```

```
[root@vz-test ~]# ifconfig eth1.105
eth1.105 Link encap:Ethernet HWaddr 00:30:48:60:71:9D
      inet addr:192.168.20.137 Bcast:192.168.20.143 Mask:255.255.255.240
      inet6 addr: fe80::230:48ff:fe60:719d/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:30 errors:0 dropped:0 overruns:0 frame:0
      TX packets:24 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:1538 (1.5 KiB) TX bytes:1596 (1.5 KiB)
```

```
[root@vz-test ~]# vzctl start 10666
Starting VE ...
VE is mounted
Adding IP address(es): 192.168.20.138
arp send: 192.168.20.138 is detected on another computer : 00:00:0c:07:ac:01
vps-net_add WARNING: arp send -c 1 -w 1 -D -e 192.168.20.138 eth1.106 FAILED
arp send: 192.168.20.138 is detected on another computer : 00:1a:6d:02:43:45
vps-net_add WARNING: arp send -c 1 -w 1 -D -e 192.168.20.138 eth1.104 FAILED
Setting CPU units: 1000
Setting CPUs: 1
Configure meminfo: 768000
File resolv.conf was modified
VE start in progress...
```

```
[root@vz-test ~]# ping -c 1 192.168.20.138
PING 192.168.20.138 (192.168.20.138) 56(84) bytes of data.
64 bytes from 192.168.20.138: icmp_seq=1 ttl=64 time=0.075 ms
```

```
--- 192.168.20.138 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
```

rtt min/avg/max/mdev = 0.075/0.075/0.075/0.000 ms

```
[root@vz-test ~]# ifconfig eth1.105
eth1.105: error fetching interface information: Device not found
```

```
[root@vz-test ~]# cat /proc/net/vlan/config
VLAN Dev name | VLAN ID
Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD
eth1.106 | 106 | eth1
eth1.104 | 104 | eth1
```

```
[root@vz-test ~]# ifconfig eth1.105
eth1.105: error fetching interface information: Device not found
```

```
[root@vz-test ~]# vzctl stop 10666
Stopping VE ...
VE was stopped
VE is unmounted
```

```
[root@vz-test ~]# ifconfig eth1.105
eth1.105 Link encap:Ethernet HWaddr 00:30:48:60:71:9D
          BROADCAST MULTICAST MTU:1500 Metric:1
          RX packets:30 errors:0 dropped:0 overruns:0 frame:0
          TX packets:22 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1538 (1.5 KiB) TX bytes:1428 (1.3 KiB)
```

```
[root@vz-test ~]# cat /proc/net/vlan/config
VLAN Dev name | VLAN ID
Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD
eth1.106 | 106 | eth1
eth1.104 | 104 | eth1
eth1.105 | 105 | eth1
```

Subject: Re: VLAN Setup
Posted by [vaverin](#) on Thu, 09 Aug 2007 09:51:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

As far as I understand you want to use several VE with vlan's IP without any vlan configurations inside VEs. Am I correct?

jonasb wrote on Thu, 09 August 2007 12:57 I can't get it to work. I can start the VE. And I can ping it from HN but not outside HN. The strange thing is that my vlan disappear when I start the VE. And is available again when I stop it.

```
[root@vz-test ~]# uname -a
Linux vz-test 2.6.18-8.1.8.el5.028stab039.1PAE #1 SMP Mon Jul 23 20:45:26 MSD 2007 i686
i686 i386 GNU/Linux
```

```
[root@vz-test ~]# vconfig add eth1 105
Added VLAN with VID == 105 to IF -:eth1:-
```

```
[root@vz-test ~]# ifconfig eth1.105 192.168.20.137 netmask 255.255.255.240 broadcast
192.168.20.143
```

here you have configured vlan vid 105 on HW node. OK

jonasb wrote on Thu, 09 August 2007 12:57

```
[root@vz-test ~]# vzctl set 10666 --netdev_add eth1.105 --save
Saved parameters for VE 10666
```

This is wrong, you should not move this device into VE, it's extra task for your purposes.

jonasb wrote on Thu, 09 August 2007 12:57

```
[root@vz-test ~]# vzctl set 10666 --ipadd 192.168.20.138 --save
Saved parameters for VE 10666
```

This is wrong again: 192.168.20.138 is already used. You should use any non-used IP for your VEs.

Thank you,
Vasily Averin

Subject: Re: VLAN Setup
Posted by [jonasb](#) on Thu, 09 Aug 2007 10:45:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

vaverin wrote on Thu, 09 August 2007 05:51As far as I understand you want to use several VE with vlan's IP without any vlan configurations inside VEs. Am I correct?

Correct!

vaverin wrote on Thu, 09 August 2007 05:51

jonasb wrote on Thu, 09 August 2007 12:57I can't get it to work. I can start the VE. And I can ping it from HN but not outside HN. The strange thing is that my vlan disappear when I start the VE. And is available again when I stop it.

```
[root@vz-test ~]# uname -a
Linux vz-test 2.6.18-8.1.8.el5.028stab039.1PAE #1 SMP Mon Jul 23 20:45:26 MSD 2007 i686
i686 i386 GNU/Linux
```

```
[root@vz-test ~]# vconfig add eth1 105
Added VLAN with VID == 105 to IF -:eth1:-
```

```
[root@vz-test ~]# ifconfig eth1.105 192.168.20.137 netmask 255.255.255.240 broadcast
192.168.20.143
```

here you have configured vlan vid 105 on HW node. OK

jonasb wrote on Thu, 09 August 2007 12:57

```
[root@vz-test ~]# vzctl set 10666 --netdev_add eth1.105 --save
Saved parameters for VE 10666
```

This is wrong, you should not move this device into VE, it's extra task for your purposes.

OK, so I don't use --netdev_add at all?

vaverin wrote on Thu, 09 August 2007 05:51

jonasb wrote on Thu, 09 August 2007 12:57

```
[root@vz-test ~]# vzctl set 10666 --ipadd 192.168.20.138 --save
Saved parameters for VE 10666
```

This is wrong again: 192.168.20.138 is already used. You should use any non-used IP for your VEs.

It doesn't matter which IP-address I use I always got error messages that they are in use. I'm really sure that I don't use them anywhere else.

Subject: Re: VLAN Setup

Posted by [vaverin](#) on Thu, 09 Aug 2007 11:09:07 GMT

[View Forum Message](#) <> [Reply to Message](#)

jonasb wrote on Thu, 09 August 2007 14:45 OK, so I don't use --netdev_add at all?

exactly, you should only assign non-used IP for your VE

jonasb wrote on Thu, 09 August 2007 14:45 It doesn't matter which IP-address I use I always got error messages that they are in use. I'm really sure that I don't use them anywhere else.

Hmm, it looks strange. you need to understand where is interface with MAC addresses

00:00:0c:07:ac:01 and 00:1a:6d:02:43:45

Could you show "arp -an" output on HW node?

Or probably you can send me access credential via PM -- I'll try to investigate this issue on your node.

Thank you,
Vasily Averin

Subject: Re: VLAN Setup
Posted by [jonasb](#) on Sat, 11 Aug 2007 05:48:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

It worked a little better with source based routing, http://wiki.openvz.org/Source_based_routing

The problem with warnings about IP already in used is still there.

After a couple of hours (3-5 hours) it stops working.
If I restart the VE or if remove and add source based routing it begin to work again.

ip route, ip address, ip rule is the same all the time.

Does anyone have a clue?

Subject: Re: VLAN Setup
Posted by [jonasb](#) on Mon, 13 Aug 2007 19:46:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

This is now solved. Andrey@SWSOft helped me out.
The problem was when the router cleaned up the arp cache the arp for the VE was removed.
I had to add an extra "ip rule" and after that it worked.

```
/sbin/ip rule add from 192.168.20.128/28 table 90  
/sbin/ip rule add to 192.168.20.128/28 lookup main pref 90  
/sbin/ip route add default dev eth1.10 via 192.168.20.130 table 90
```

Subject: Re: VLAN Setup
Posted by [jonasb](#) on Fri, 25 Jan 2008 08:17:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

This is not solved yet.

I use the following solution to get several VLANs to work on the same HN.
It works.

```
echo 1 > /proc/sys/net/ipv4/conf/eth1.100/proxy_arp
echo 0 > /proc/sys/net/ipv4/conf/eth1.100/rp_filter
```

But when you have 2 (or more) HNs in the VLAN you have problem.

Since I disabled rp_filter and enabled proxy_arp both HNs try to answer the incoming ARP request.

And the HN who can answer fastest will be arp cached in the router.

So if the wrong HN answers that MAC address will be used in the router and the packets will not get through.

When I manually delete the ARP entry and add the right one it works.

Any ideas how to solve this? There must be some way to get real VLAN support in OpenVZ.

Subject: Re: VLAN Setup
Posted by [jonasb](#) on Mon, 28 Jan 2008 20:58:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

Any ideas? Should I make a bug report?
